AFRL-IF-RS-TR-1998-140
Final Technical Report
July 1998

ASYNCHRONOUS, DISTRIBUTED, SCALABLE ALGORITHMS FOR INTELLIGENT REASONING WITH GEOGRAPHICALLY DISPERSED, HYBRID KNOWLEDGE BASES

University of Maryland

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. 5D30

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-1998-140 has been reviewed and is approved for publication.

APPROVED: Joseph A. Carozzoni
JOSEPH A. CAROZZONI
Project Engineer

FOR THE DIRECTOR: Northrup Fowler
NORTHRUP FOWLER, III, Technical Advisor
Information Technology Division
Information Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/IFTB, 525 Brooks Road, Rome, NY 13441-4505. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

ASYNCHRONOUS, DISTRIBUTED, SCALABLE ALGORITHMS FOR INTELLIGENT REASONING WITH GEOGRAPHICALLY DISPERSED, HYBRID KNOWLEDGE BASES

V. S. Subrahmanian

Contractor: University of Maryland
Contract Number: F30602-93-C-0241
Effective Date of Contract: 15 September 1993
Contract Expiration Date: 14 September 1997
Short Title of Work: Asynchronous Distributed, Scalable, Algorithms for Intelligent Reasoning with Geographically Dispersed, Hybrid Knowledge Bases

Period of Work Covered: Sep 93 - Sep 97

Principal Investigator: V. S. Subrahmanian
Phone: (401) 405-6707
AFRL Project Engineer: Joseph Carozzoni
Phone: (315) 330-7796

This research was supported by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by Joseph Carozzoni, AFRL/IFTB, 525 Brooks Road, Rome, NY 13441-4505.

REPORT DOCUMENTATION PAGE			*Form Approved* *OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	July 1998	Final, Sep 93 - Sep 97

4. TITLE AND SUBTITLE

ASYNCHRONOUS, DISTRIBUTED, SCALABLE ALGORITHMS FOR INTELLIGENT REASONING WITH GEOGRAPHICALLY DISPERSED, HYBRID KNOWLEDGE BASES

5. FUNDING NUMBERS

C - F30602-93-C-0241
PE - 62301E
PR - A716
TA - 00
WU - 01

6. AUTHOR(S)

V. S. Subrahmanian

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Maryland
Office of Research Administration and Advancement
College Park MD 20742

8. PERFORMING ORGANIZATION REPORT NUMBER

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington VA 22203-1714

Air Force Research Laboratory/IFTB
525 Brooks Road
Rome NY 13441-4505

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

AFRL-IF-RS-TR-1998-140

11. SUPPLEMENTARY NOTES

Air Force Research Laboratory Project Engineer: Joseph A. Carozzoni/IFTB/(315) 330-7796

12a. DISTRIBUTION AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

Integrating data and knowledge from multiple heterogeneous sources, each one possibly with a different underlying data model, is not only an important aspect of automated reasoning, but also of retrieval systems where queries can span such multiple sources. These sources can be as different as relational or deductive databases, object bases, constrained data (e.g. knowledge bases), structured files (e.g. spreadsheets), or even arbitrary program packages encapsulating specific knowledge, often in a hard-wired form accessible only through function calls. Many queries can only be answered if data and knowledge from these different sources are available. In 1991-92, Gio Wiederhold proposed the pioneering concept of a mediator - a program that integrates multiple databases. The principal goal of this project was to develop a platform for the creation of mediated applications. Such a platform would provide a mechanism within which mediators may be developed for a variety of applications. The platform itself would be application independent, and would provide a variety of underlying technology and services that would be critical to the success of any specific application involving the use of mediation technology. The product of this research is a system called WebHERMES (Heterogeneous Reasoning and Mediation System).

14. SUBJECT TERMS	15. NUMBER OF PAGES
Software, Knowledge-Based Systems, Intelligent Database	36
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

1 Project Goals

Integrating data and knowledge from multiple heterogeneous sources, each one possibly with a different underlying data model, is not only an important aspect of automated reasoning, but also of retrieval systems where queries can span multiple such sources. These sources can be as different as relational or deductive databases, object bases, (constraint) knowledge bases, or even (structured) files and arbitrary program packages encapsulating specific knowledge, often in a hard-wired form accessible only through function calls. Many queries can only be answered if data and knowledge from these different sources are available.

In 1991–92, Gio Wiederhold proposed the pioneering concept of a *mediator* – a program that integrates multiple databases. However, while the goals of precisely what objectives a mediator would satisfy were clear, how these objectives would be accomplished and implemented was not clear. *The principal goal of this project was to develop a platform for the creation of mediated application. Such a platform would provide a mechanism within which mediators may be developed for a variety of application. The platform itself would be application independent, but would provide a variety of underlying technology and services that would be critical to the success of any specific application involving the use of mediation technology.*

In this project, we have developed a formal, theoretically solid framework for the creation and deployment of mediators that access distributed data sources, and shown that this mathematically justified framework scales up to large scale applications involving integrated access not only to multiple databases, but also to multiple

[1]E-mail: vs@cs.umd.edu, Phone: (301) 405-2711, FAX: (301) 405-6707.

data structures, and software packages located at diverse networked sources. The resulting system, called WebHERMES, is accessible to any user who has access to the world-wide web through any standard Web browser. This includes access from Unix workstations, PC's, MACs, as well as palmtop computing devices such as the Philips Velo or the US Robotics Pilot.

The organization of this report is as follows. Section II explains the main scientific contributions of this project. Section III explains the software that has been developed. Section IV specifies what Educational Objectives have been accomplished from this project. Section V presents a list of all publications acknowledging this contract.

2 Awards/Recognition

The HERMES project, and its participants, have received significant recognition for their work on this project, from a variety of external sources. These are listed below:

- **National Young Investigator Award**, 1993 to V.S. Subrahmanian (PI), National Science Foundation.
- **Maryland Distinguished Young Scientist Award**, V.S. Subrahmanian (PI). Maryland Science Center and the Maryland Academy of Sciences, 1997.
- **Association for Computing Machinery (ACM) Washing Chapter Samuel Alexander Award**, 1997, Kasim S. Candan (graduate research assistant funded by this contract), for an outstanding doctoral dissertation.
- **Business Week Magazine** highlights the accomplishments of Sibel Adali who received her PhD for her work on this project.
- **Publications:** Over 30 publications in top-quality, archival scientific journals, and 16 papers in leading scientific conferences were published due, in part, to support received under this contract.

3 Scientific Accomplishments of Project

The HERMES (Heterogeneous Reasoning and Mediator System) project was started in Sep. 1993. During this time, we have developed:

- A language in which mediators can be expressed
- A compiler within which mediators expressed in the above language can be implemented
- A distributed computation framework so that the mediator compiler can access data at multiple sites across the network

- A set of techniques to optimize queries to such distributed heterogeneous repositories
- A set of techniques to incrementally create materialized mediated views (better known as data warehouses) consisting of information from multiple sources
- A set of techniques to specify security policies in mediated systems, as well as process updates in secure mediators
- Web client access to mediated applications
- A unified framework for representing and manipulating multimedia data located across the Internet.

We will now describe briefly, our contributions in each of these areas.

3.1 Mediator Language

We have proposed the following concepts for the HERMES mediator language. A domain, D, is an abstraction of databases and software packages and consists of three components: (1) a set, S, whose elements may be thought of as the data-objects that are being manipulated by the package in question, (2) a set F of functions on S – these functions take objects in S as input, and return, as output, objects from their range (which needs to be specified). The functions in F may be thought of as the predefined functions that have been implemented in the software package being considered, (3) a set of relations on the data-objects in S – intuitively, these relations may be thought of as the predefined relations in the domain,

In our system, called HERMES ("Heterogeneous Reasoning and Mediator System"), a domain call is a syntactic expression of the form

$$\texttt{domainname} : \texttt{domainfunction}(< \texttt{argument1}, ..., \texttt{argumentn} >)$$

where `domainfunction` is the name of the function. and `argument1,...,argumentn` are arguments to that function. Intuitively, a domain call may be read as: in the domain called domainname, execute the function called domainfunction on the arguments

$$< \texttt{argument1}, ..., \texttt{argumentn} > .$$

The result of executing this domain call is coerced into a set of entities that have the same type as the output type of the function domainfunction on the arguments

$$< \texttt{argument1}, ..., \texttt{argumentn} > .$$

A domain-call atom DCA-atom) is of the form

$$\texttt{in}(\texttt{X}, \texttt{domainname} : \texttt{domainfunction}(< \texttt{arg1}, \ldots, \texttt{argn} >))$$

polymorphic set membership predicate. For example,

$$in(A, paradox : select_eq('phonebook', "name", "josmith"))$$

is a DCA-atom that is true just in case A is a tuple in the result of executing a selection operation (finding tuples where the NAME field is JO SMITH on a relation called PHONEBOOK maintained in a PARADOX database system.

A mediator is a set of rules of the form

$$A \leftarrow D1\&\ldots\&Dm|A1\&..\&An$$

where $A1, ..., An$ are atoms. and $D1, ..., Dm$ are DCA-atoms.

We have studied the syntax and semantics of this language exhaustively, yielding a clean amalgamation of multiple databases. data structures and software packages. We have developed algorithms that are provably correct that answer queries to these databases very efficiently.

4 Mediator Compiler

We have built a mediator compiler within which queries to HERMES mediators may be expressed and processed. There are two important aspects to constructing a mediator: domain integration and semantic integration. Intuitively, domain integration is the physical linking of the data sources and reasoning systems, while semantic integration is the coherent extraction and combination of the information provided by the data and reasoning sources, serving a given purpose.

The HERMES compiler takes as input. a mediator expressed in the HERMES language expressed in the preceding section, and produces as output, a set of data structures that may be used to process and execute queries in the Hermes query language. When a user of an application mediator built in HERMES expresses a query, the mediator rules defined with the application mediator expands the query into a set of subqueries. Such subqueries may be subqueries either to the HERMES system itself, or to external data sources accessed by the HERMES mediator. Here is an example of how a HERMES query is processed.

Example Query: Let **rte1** be a ternary predicate such that $\mathtt{rte1}(O, D, R)$ is satisfied iff R is a route from the origin to an unspecified destination such that the destination has an airfield as well as certain types of ammunition. For this, we may define the following clause in the mediator:

```
rte1(O,D,R) ←
        in(P1,paradox : select=(facilities,facility,"airfield")) &
        in(P2,dbase : select=(supplies.item,"ammunition")) &
        = (P1.place,P2.place) &
        in(D,spatial : findpt(P1.place)) &
        in(R,rp : route(O,D)).
```

To obtained the result from a given location ℓ, we can pose the query

$$\leftarrow \mathtt{rte1}(\ell, \mathtt{D}, \mathtt{R}).$$

This is then processed as follows: `PARADOX` is invoked which `SELECT`s all tuples from the `facilities` relation that have the an `airfield` facility. `P1` is then instantiated to one of the selected tuples. Next `DBASE` is then queried to `SELECT` all tuples from the `supplies` relation that have the `item` field set to `ammunition`. `P2` is instantiated to one such tuple. A check is made to see that `P1` and `P2` have the same `place` field. In other words, this ensures that *single* place is found with both ammunition and an airfield? If not, the HERMES inference engine looks for other possible instantiations of `P1` and `P2` that satisfy these constraints. Finally, the xy-location of the place `P1.place` is computed using the spatial domain, instantiating `D`, and `RP` is called to find a route from the origin ℓ to D.

The HERMES query processing algorithm is a sound and complete algorithm for processing queries to heterogeneous mediated systems.

5 Query Optimization and Caching

An important issue that we have studied is ways to make the processing of queries in heterogeneous reasoning systems more efficient. We advocate the intelligent use of high-speed caches to avoid computations whenever possible. To accomplish this we introduce the concept of an "invariant", i.e. an expression about the known input/output relationships of a program that can be processed by the mediator. We have shown how such caches may be maintained, and how the query processing procedure can make better use of these caches, given the knowledge about different packages, to reduce the complexity of query execution. Our methods are sound and complete.

It is possible that in the processing of the rules in the mediator, "similar" function calls to external programs will need to be executed several times since the same kind of information may be requested over and over by different users. Backtracking is another reason for such a situation. Calling an external program is usually a costly operation because of the memory, CPU requirements and possible network delays. Furthermore, actual packages may levy charges for accessing them. Suppose there is a way to guess "some" of the answers that will be returned by an external call. If a refutation is found by substituting one of these answers, then there is no need to execute the external domain call. This is accomplished by caching the answers returned by previous external calls and re-using them when needed. Similarly, if there is a way of knowing that a function is not defined for some inputs, whenever it is called for these inputs, we can terminate the search down a path of the search space.

The challenge of this approach lies in deciding which sets of answers are relevant, and in representing the input/output behavior of some external functions. This information is stored in the system with the help of some explicit rules which will be referred to as "invariants". Invariants are expressions specifying the relation between the set of answers returned by an external call, its arguments and other possible external calls.

As an example, consider the following invariant :

$$\texttt{T2} \geq \texttt{T1} \Rightarrow \texttt{f(T2)} \supseteq \texttt{f(T1)}$$

The above expression can be read as follows: if $\texttt{T2} \geq \texttt{T1}$, then all the solutions of `f(T1)` are also solutions of `f(T2)`. Hence, if the set of answers for `f(T1)` were previously stored in a cache, then these answers can be re-used whenever the function `f(T2)` is called; if none of the answers to `f(T1)` satisfies the rest of the query then `f(T2)` needs to be computed. An example of such an invariant is given below:

Example 1 Suppose **relation** is a constant in the mediator which refers to a relational database with the usual selection operators. For example, $\texttt{select}_{\leq}\texttt{(R,F,V)}$ selects all the tuples in table `R` such that the value of the field `F` is less than or equal to `V`. Then, the following are possible invariants for different **select** functions.

$$\begin{aligned}
\texttt{T2} \leq \texttt{T1} \Rightarrow\ & \texttt{relation}:\texttt{select}_{\leq}\texttt{(R,Field,T1)} \supseteq \\
& \texttt{relation}:\texttt{select}_{\leq}\texttt{(R,Field,T2)}. \\
\texttt{T2} \geq \texttt{T1} \Rightarrow\ & \texttt{relation}:\texttt{select}_{\geq}\texttt{(R,Field,T2)} \supseteq \\
& \texttt{relation}:\texttt{select}_{\geq}\texttt{(R,Field,T1)}.
\end{aligned}$$

The first invariant can be read as: For any given database `R` and field **Field** in the domain **relation**, whenever $\texttt{T2} \geq \texttt{T1}$ is satisfied, all the tuples that are in $\texttt{relation}:\texttt{select}_{\geq}\texttt{(R,Field,T2)}$, are also in $\texttt{relation}:\texttt{select}_{\geq}\texttt{(R,Field,T1)}$.

Example 2 Suppose the domain **spatial** is a spatial data structure such as a point quadtree storing points in two-dimensional space. The function **vertical_slice(File,X,Dist)** in this domain returns all the points that have X-coordinates between **X+Dist** and **X-Dist**, in other words all the points that are in the vertical slice taken from **X-Dist** to **X+Dist**. The following is an invariant about this function:

$$\begin{aligned}
&\texttt{Dist1} \leq \texttt{Dist2} \Rightarrow \\
&\quad \texttt{spatial}:\texttt{vertical_slice(File,X,Dist2)} \supseteq \\
&\quad \texttt{spatial}:\texttt{vertical_slice(File,X,Dist1)}.
\end{aligned}$$

which states that whenever the X-coordinate is fixed, the points in a vertical slice are contained in any of the bigger vertical slices. We can easily write similar invariants for other spatial functions. The invariant for the **horizontal_slice** function is the same as **vertical_slice**. As for the **range(X,Y,Dist)** function which returns all the points that are at distance **Dist** from point **(X,Y)** (i.e. all points **(X1,Y1)** such that $(\texttt{X}-\texttt{X1})^2+(\texttt{Y}-\texttt{Y1})^2 \leq \texttt{Dist}^2$ we can write the following invariants:

$$\begin{aligned}
&\texttt{Dist1} \leq \texttt{Dist2} \Rightarrow \\
&\quad \texttt{spatial}:\texttt{range(File,X,Y,Dist2)} \supseteq \\
&\quad \texttt{spatial:range(File,X,Y,Dist1)}. \\
&|\texttt{X1}-\texttt{X2}| \leq |\texttt{Dist1}-\texttt{Dist2}| \Rightarrow \\
&\quad \texttt{spatial}:\texttt{range(File,X2,Y,Dist2)} \supseteq \\
&\quad \texttt{spatial:range(File,X1,Y,Dist1)}.
\end{aligned}$$

Our research proposes a cost-based optimization technique that caches statistics of actual calls to the sources and consequently estimates the cost of the possible execution plans based on the statistics cache. We investigate issues pertaining We investigate issues pertaining to the design of the statistics cache and experimentally analyze various tradeoffs. We also present a query result caching mechanism that allows us to effectively use results of prior queries when the source is not readily available. We employ the novel *invariants* mechanism, which shows how semantic information about data sources may be used to discover cached query results of interest.

6 Maintaining Mediated Views/Warehouses

A *mediated materialized view* (often called a data warehouse) is a view of a body of distributed heterogeneous data that is precomputed and stored as a cache of the sort described in the preceding section. As in the case of traditional views, mediated views are materialized for efficiency reasons. A materialized view can be affected by two kinds of updates, namely updates to the materialized view, and updates to the underlying sources.

If an update of the first kind occurs to a view, whether materialized or not, the problem of reflecting the update correctly by changing the base tables appropriately needs to be addressed. This problem is called the *view update problem* and has been discussed extensively for relational, deductive, and object-oriented databases. However, our objective is slightly different. We do not necessarily assume that an update occurring to a view has to be reflected within some underlying source. Instead, we assume that the view itself — or, to be more precise, its definition — is affected by the update. This kind of update affecting the view's definition is typically not treated within the view update literature. One exception are deductive databases, where the addition or deletion of rules to the definition of an intensional predicate is discussed by Teniente. However, they neither materialize nor preprocess the view for efficiency reasons.

Within the traditional context, the second case occurs if an update to a base table occurs which possibly affects a materialized view. The resulting problem — preserving the consistency of the view — is called *view maintenance*. However, since we do not necessarily materialize the view upon the underlying sources of our mediated views but instead perform materialization by unfolding the view definition as independent as possible from the underlying sources, the traditional view maintenance problem occurs quite differently to us. Hence, the traditional view maintenance problem and our problem do not intersect but complement each other.

Subsequently, we treat both kinds of updates to materialized mediated views and show how they can be handled efficiently. More specifically, the primary aim is to specify how to efficiently maintain views of mediated systems such as those that may be constructed in HERMES when insertion and deletion requests of both of the above two kinds are made. As in the standard case, a *materialized view* in mediated systems may be thought of as a set of facts that can be concluded from the mediator rules.

However, we show that more generally, a materialized mediated view may be regarded as a set of *constraint* atoms that are not necessarily ground. Taking materialized views to be sets of constrained atoms leads to a number of advantages:

1. First of all, it allows us to perform updates to *constrained databases* as well as mediated systems. To our knowledge, there are currently no methods to incrementally maintain views in constrained databases.

2. We show for updates of the second kind that even in the case of unconstrained databases, such as those considered by Gupta, Mumick and Subrahmanian, (which we have been told is now used by AT& T for billing purposes) this approach leads to a simpler and more efficient deletion algorithm than the deletion algorithm, DRed presented in earlier.

In other words, not only have we developed efficient algorithms for view management, these algorithms also (in some cases) improve upon existing algorithms for view management in traditional relational databases.

7 Ontology Management

Any mediator, in integrating heterogeneous sources, has to resolve both syntactic and semantic conflicts between (the data in) in the disparate sources. While considerable work has been done on this problem in the context of multi-database systems, little algorithmic support has been developed for resolving (especially semantic) conflicts, and currently, resolving them is largely a responsibility of the mediator developer. We develop appropriate concepts and algorithms for solving the following problems.

- Resolving the conflicts between data coming from heterogeneous sources

- Allowing users to personalize queries so as to address their own needs (e.g., a user from India might want prices returned in Indian Rupees).

- Answering personalized queries.

- Maintaining a mediator against changes to the data sources, in the form of restructuring. Such restructuring may be motivated by the requirements of the local user community of the source. Our ideas and techniques apply to any mediator framework, such as the TSIMMIS project at Stanford University, the HERMES project at University of Maryland, the SchemaLog project at Concordia University, the Disco project at Inria, and several others.

8 Security in Mediated Systems

Over the last few years, there has been considerable work on security in databases. ost of this work has been limited to the realm of relational databases though of late,

some work has been done on extending these security paradigms to object-oriented databases deductive databases, and other paradigms. Castano et. al. provide a comprehensive description of related work. Despite the differences in the underlying data paradigm, all these frameworks share a single trait that we (cynically?) term the principle of paranoia.

> THE PRINCIPLE OF PARANOIA. The DBMS must take all steps necessary in order to insure that the user u cannot to infer any item in a pre-desig‑ated set $\mathcal{S}(u)$ of items that are to be kept secret from the user.

However, with the evolution of the information superhighway, there is now an immense amount of information available in a very wide variety of databases. Wiederhold has proposed the concept of a *mediator* – intuitively, a mediator is a program that integrates multiple databases. Consider a mediator program M that integrates some software packages $P_1, \ldots, P_k$. Each of the packages $P_1, \ldots, P_k$ may enforce its own unique *local* security policy. Some may represent completely "open-source" software/data, while others may place certain restrictions on the use of certain facilities and/or data residing within it. In contrast to the principle of paranoia commonly enforced in ordinary databases, mediated systems must attempt to be *maximally cooperative* to the user, yet at the same time, they must respect the security constraints of the individual databases/packages participating in the mediated system. Thus, for instance, two packages P_1 and P_2 may both be able to satisfy a user's request – however, package P_1 uses secure data, while package P_2 uses open-source data. In this case, the mediated system may reasonably use package P_2 to respond to the user's query, even though package P_1 feels this data should be kept hidden from the user. Notice that in this case, the user could directly query P_2 and get the data without using the mediator at all, so the mediator might as well do it for him, unless a *global* security condition maintained by the mediator prevents this. Thus, in the case of mediated systems, we may wish to implement a policy of *cautious cooperation.*

> THE PRINCIPLE OF CAUTIOUS COOPERATION. If a user's query can be answered using open-source information, then the mediator will answer the query unless doing so will *directly* violate the system's *global* security constraints. However, the system will always respect the rights of individual packages participating in the mediated system, and ensure that no single package violates its own local security policy.

The principle of cautious cooperation ensures that a given query will cause no direct violation of global integrity constraints, but may leave the path open for future violations or for inferential violations. A slightly more conservative policy, that we term the principle of *conservative cautious cooperation*, will answer a query posed by the user iff the answer to that query will not yield an "inference path" (sequence of open-source queries coupled with logical reasoning) that the user may use to violate security (unless such a path existed prior to the query being issued by the user).

THE PRINCIPLE OF CONSERVATIVE CAUTIOUS COOPERATION. If a user's query can be answered using open-source information. then the mediator will answer the query unless doing so will cause there to be a sequence of queries (that only reflect open-source accesses) such that if the user asks this sequence of queries. then he will be able to violate the system's *global* security constraints. However. the system will always respect the rights of individual packages participating in the mediated system. and ensure that no single package violates its own local security policy.

In our research. we have developed techniques by which mediators may be efficiently and scalably extended to encode the principles of paranoia and the principle of cautious cooperation. as well as the principle of conservative cautious cooperation.

9 Web Access to Mediator Technology

We have developed algorithms which will take as input, a HERMES mediator M, and generate as output. an HTML-form that can be used to query mediator M by any user who has access to the World Wide Web through standard Web browsers such as Netscape and Microsoft Explorer. In particular. as a consequence of this Web-page generation module. users with Web browsers on:

- Unix workstations (e.g. SUNs and DECs)
- PC devices (e.g. IBM-PC compatibles)
- Wireless/cellular palmtop devices (e.g. The Philips Velo. US Robotics Pilot)

can now access WebHERMES from such devices.

10 Heterogeneous Multimedia Databases

Though numerous multimedia systems exist in the commercial market today, relatively little work has been done on developing the mathematical foundations of multimedia technology. We attempt to take some initial steps towards the development of a theoretical basis for multimedia information system. To do so, we develop the notion of a structured multimedia database system. We begin by defining a mathematical model of a media-instance. A media-instance may be thought of as "glue" residing on top of a specific physical media-representation (such as video, audio, documents, etc.) Using this "glue", it is possible to define a general purpose logical query language to query multimedia data. This glue consists of a set of "states" (e.g. video frames, audio tracks, etc.) and "features", together with relationships between states and/or features. A structured multimedia database system imposes a certain mathematical structure on the set of features/states. Using this notion of a structure, we are able to define indexing structures for processing queries, methods to relax queries when answers do not exist to those queries, as well as sound, complete and terminating

procedures to answer such queries (and their relaxations, when appropriate). We show how a media-presentation can be generated by processing a sequence of queries, and furthermore we show when these queries are extended to include *constraints*, then these queries can not only generate presentations, but also generate temporal synchronization properties and spatial layout properties for such presentations. We describe the architecture of a prototype multimedia database system based on these principles.

11 An Application to Terrain Reasoning

The work described here was done jointly with researchers at the US Army Topographic and Engineering Center in Ft. Belvoir, VA.

In this section, we will describe an application of our work to intelligent terrain reasoning that involves integrating terrain map data, relational data, and planning packages (developed at the US Army Corps of Engineers). The purpose of such an integrated system is many-fold. It can be used as a basis for vehicular navigation in disaster relief situations (e.g. floods, earthquakes, volcanic disasters, etc.), as well as in military mission planning applications. In these applications, a user, who may either be a human or may be an autonomous vehicle, may be interested in posing queries of the following types:

- **(Unknown Destination)** Given a location, find a place that has an airfield as well as certain types of ammunition. Presumably these resources are needed in order for the autonomous/manned vehicle to satisfy its mission.
- **(Route Properties)** Furthermore, no point in the route may be less than 4 miles from an enemy outpost. In this example, in addition to the fact that the destination is unknown, we have the fact that the query asks not only for a route to this unknown destination point, but it asks for a *route that satisfies certain desiderata*, i.e. which satisfies certain conditions that require accessing external databases (e.g. to figure out where enemy outposts lie).

A route planner (which we will call `RP`) has been implemented at the US Army Topographic and Engineering Center. Given two points, this route planner will find an optimal, least-cost path between these two points (if one exists). Thus, for instance, the query

$$\texttt{rp} : \texttt{route}((35,70),(200,98))$$

returns the set of least-cost paths from the origin point, (35,70) to the destination point (200,98) that are found by the Army's route planner.

We illustrate how this example may be solved within the HERMES framework, using `RP` as a domain. For this, let us suppose that we have a relational (`PARADOX`) database containing a relation called `facilities` having the schema (Name,Facility). Thus, this relation may contain a tuple of the form (*awasa*, *airport*) denoting that the place, Awasa, has an airport. Other tuples in the relation `facilities` may be similarly interpreted. Suppose there is another (`DBASE`) database containing a relation

called `supplies` having the schema (Place.Item) - an example tuple in this relation is $(awasa, gas)$ specifying that gas is available at Awasa[2].

Example Query:

$$\texttt{rte2(O,D,R)} \leftarrow \texttt{rte1(O,D,R)} \,\&\, \texttt{good(R)}.$$
$$\texttt{good(nil)} \leftarrow$$
$$\texttt{good(cons(H,T))} \leftarrow \texttt{goodpoint(H)} \,\&\, \texttt{good(T)}.$$
$$\texttt{goodpoint(H)} \leftarrow \texttt{is(\{\},spatial} : \texttt{range(H,4))}.$$

Note the use of the special HERMES predicate `rte1` has been described earlier.

12 Software Developed

Appendix A contains a complete user manual of the HERMES software.

13 Educational Accomplishments

During the pursuit of this research, we have accomplished the following milestones:

- V.S. Subrahmanian (PI) received the NSF National Young Investigator Award.
- V.S. Subrahmanian (PI) received the Maryland Distinguished Young Scientist Award (Maryland Academy of Sciences).
- 2 PhD's were granted:
  - Sibel Adali received her PhD in 1996 and is currently an Assistant Professor of Computer Science at Rensselaer Polytechnic Institute in Troy, NY.
  - Kasim S. Candan received his PhD in 1997 and is currently an Assistant Professor of Computer Science at Arizona State University in Tempe, AZ.
- Kasim S. Candan (student supported by this contract), received the 1997 ACM Samuel Alexander Award for an outstanding dissertation.
- Sibel Adali's work supporting this contract received Press Highlights in *Business Week Magazine*, June 23, 1997.
- The following students supported in part by this contract received Master's degrees: Charlie Ward, Vadim kagan.

[2]In practice, these relations will contain much more detail, but we keep them simple here in order to facilitate presentation.

14 Publications

The following papers were supported in part by this contract.

JOURNAL PAPERS

1. V. Kagan. A. Nerode and V.S. Subrahmanian. Computing Definite Logic Programs by Partial Instantiation, ANNALS OF PURE AND APPLIED LOGIC, Vol. 67, pps 161–182. 1994.

2. V.S. Subrahmanian. Amalgamating Knowledge Bases, ACM TRANSACTIONS ON DATABASE SYSTEMS, 19, 2, pp. 291–331, 1994.

3. C. Bell, A. Nerode. R. Ng and V.S. Subrahmanian. Mixed Integer Programming Methods for Computing Non-Monotonic Deductive Databases, JOURNAL OF THE ACM. Vol. 41. Nr. 6, pps 1178–1215, Nov. 1994.

4. J. Grant and V.S. Subrahmanian. Reasoning About Inconsistent Knowledge Bases. IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, 7,1, pps 177–189.

5. A. Nerode, R. Ng and V.S. Subrahmanian. Computing Circumscriptive Databases, Part I: Theory and Algorithms. INFORMATION AND COMPUTATION, Vol. 116, Nr. 1, pps 58–80.

6. S. Adali and V.S. Subrahmanian. Amalgamating Knowledge Bases, II: Distributed Mediators. INTL. JOURNAL OF INTELLIGENT COOPERATIVE INFORMATION SYSTEMS, 3,4. pps 349–383. 1994.

7. C. Baral, S. Kraus, J. Minker, and V.S. Subrahmanian. Combining Default Logic Databases. in: INTL. JOURNAL OF INTELLIGENT AND COOPERATIVE INFO. SYSTEMS, 3,3, pps 319–348, 1994.

8. S. Kraus and V.S. Subrahmanian. Multiagent Reasoning with Probability, Time and Beliefs, INTERNATIONAL JOURNAL OF INTELLIGENT SYSTEMS, 10(5) 459–499.

9. A. Brink, S. Marcus and V.S. Subrahmanian. Heterogeneous Multimedia Reasoning. IEEE COMPUTER, 28, 9, pps 33–39. Sep. 1995.

10. X. Y. Wang, S. Jajodia, and V.S. Subrahmanian. Temporal Modules: An Approach Toward Federated Temporal Databases. INFORMATION SCIENCES, Vol. 82, pps 103–128, 1995.

11. V.S. Subrahmanian, D.S. Nau and C. Vago. WFS + Branch and Bound = Stable Models, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, 7, 3. pps 362–377, June 1995.

12. P. Bonatti, S. Kraus and V.S. Subrahmanian. Foundations of Secure Deductive Databases, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, 7,3, pps., June 1995.

13. K. Erol, D.S. Nau and V.S. Subrahmanian. Complexity, Decidability and Undecidability Results for Domain-Independent Planning. ARTIFICIAL INTELLIGENCE Journal, 76.1-2, pps 75–88, 1995.

14. J. Grant and V.S. Subrahmanian. The Optimistic and Cautious Semantics for Inconsistent Knowledge Bases, ACTA CYBERNETICA, Vol. 12, No.1,1995.

15. J. Benton, S.S. Iyengar, W. Deng, N. Brener, and V.S. Subrahmanian. Tactical Route Planning, New Algorithms for Decomposing the Map, INTL. JOURNAL OF TOOLS IN ARTIFICIAL INTELLIGENCE, Vol. 5,Nos. 1 & 2 (1996) 199-218.

16. E. Hwang and V.S. Subrahmanian. Querying Video Libraries. JOURNAL OF VISUAL COMMUNICATION AND IMAGE REPRESENTATION, Vol. 7, No. 1, March, pp.44-60, 1996.

17. S. Marcus and V.S. Subrahmanian. Foundations of Multimedia Database Systems, JOURNAL OF THE ACM, Vol. 43, 3, pps 474–523, 1996.

18. C. Bell, A. Nerode, R. Ng and V.S. Subrahmanian. Implementing Deductive Databases by Mixed Integer Programming, ACM TRANSACTIONS ON DATABASE SYSTEMS, vol. 21, nr. 2, pps 238–269, 1996.

19. V. Kagan, A. Nerode and V.S. Subrahmanian. Computing Minimal Models by Partial Instantiation. THEORETICAL COMPUTER SCIENCE, vol. 155, pps 157–177.

20. S. Adali and V.S. Subrahmanian. Amalgamating Knowledge Bases, III: Algorithms, Data Structures and Query Processing. JOURNAL OF LOGIC PROGRAMMING, Vol 28(1), pps 57-100, July 1996.

21. S. Adali, K.S. Candan, S.-S. Chen, K. Erol, and V.S.Subrahmanian. Advanced Video Information Systems, ACM MULTIMEDIA SYSTEMS JOURNAL, 4, pps 172–186, 1996.

22. J. Lu, A. Nerode and V.S. Subrahmanian. Hybrid Knowledge Bases, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, 8, 5, pps 773–785, Oct. 1996.

23. A. Nerode, J. Remmel and V.S. Subrahmanian. Annotated Nonmonotone Rule Systems, THEORETICAL COMPUTER SCIENCE, Volume 171, issue 1-2, pages 77–109, Jan. 1997.

24. G. Gottlob, S. Marcus, A. Nerode, G. Salzer and V.S. Subrahmanian. A Non-Ground Realization of the Stable and Well-Founded Semantics. THEORETICAL COMPUTER SCIENCE, vol. 166, Nr. 1/2, pps 221-262, Oct.1996.

25. S. Pradhan, J. Minker and V.S. Subrahmanian. Combining Databases with Prioritized Information. accepted for publication in: JOURNAL OF INTELLIGENT INFORMATION SYSTEMS.

26. A. Brogi, V.S.Subrahmanian and C. Zaniolo. The Logic of Total and Partial Order Plans: A Deductive Database Approach. accepted for publication in: ANNALS OF MATH AND ARTIFICIAL INTELLIGENCE.

27. K.S. Candan, J. Grant and V.S. Subrahmanian. A Unified Treatment of Null Values using Constraints. Accepted for publication in: INFORMATION SCIENCE journal.

28. P. Bonatti, M.-L. Sapino and V.S. Subrahmanian. (1996) *Merging Heterogeneous Security Orderings*, accepted for publicatiopn in: JOURNAL OF COMPUTER SECURITY.

29. V.S. Lakshmanan, N. Leone, R. Ross and V.S. Subrahmanian. ProbView: A Flexible Probabilistic Database System. accepted for publication in: ACM TRANSACTIONS ON DATABASE SYSTEMS, May 1996.

30. K.S. Candan, V.S. Subrahmanian and P. Venkat Rangan. Collaborative Multimedia Systems: Synthesis of Media Objects. accepted for publicatiopn in: IEEE TRANS. ON KNOWLEDGE AND DATA ENGINEERING.

31. K.S. Candan, E. Hwang and V.S. Subrahmanian. An Event-Based Model for Continuous Media Data on Heterogeneous Disk Servers. accepted for publicatiopn in: ACM Multimedia Systems Journal.

32. K.S. Candan, B. Prabhakaran and V.S. Subrahmanian. Retrieval Schedules Based on Resource Availability and Flexible Presentation Specifications. accepted for publicatiopn in: ACM Multimedia Systems Journal.

CONFERENCE PAPERS

33. J. Benton and V.S. Subrahmanian. Using Hybrid Knowledge Bases for Missile Siting Problems. Proc. 1994 Conf. on Artificial Intelligence Applications, pps 141–148, IEEE Computer Society.

34. J. Lu, G. Moerkotte, J. Schue, and V.S. Subrahmanian. Efficient Maintenance of Materialized Mediated Views, accepted for publication in: *Proc. 1995 ACM SIGMOD Conf. on Management of Data*, San Jose, CA, May 1995.

35. V.S. Subrahmanian and C. Zaniolo. Relating Stable Models and AI Planning Domains, accepted for publication in: *Proc. 1995 Intl. Conf. on Logic Programming*, Tokyo, Japan.

36. Sibel Adali and V.S. Subrahmanian. Intelligent Caching in Hybrid Knowledge Bases, in: *Proc. 1995 Intl. Conf. on Very Large Knowledge Bases* (ed. N. Mars), pps 247–256, IOS Press, Twente, The Netherlands, May 1995.

37. The PARK Semantics for Active Rules, accepted for publication in: *Proc. 1996 Intl. Conf. on Extending Database Technology*. Avignon, France, March 1996. (with G. Gottlob and G. Moerkotte).

38. Secure Mediated Databases, in: *Proc. 1996 Intl. Conf. on Data Engineering*, Feb. 1996, New Orleans, LA. (with K.S. Candan and S. Jajodia).

39. Query Processing in Distributed Mediated Systems, in: *Proc. 1996 ACM SIGMOD Conf. on Management of Data*, Montreal, Canada, June 1996. (with S. Adali, K.S.Candan and Y. Papakonstantinou).

40. Towards a Theory of Collaborative Multimedia, Proc. 1996 IEEE Multimedia Systems Conference, pps 279–283, Hiroshima, Japan, June 1996.

41. A Deductive Database Approach to Planning in Uncertain Environments, Proc. 1996 Intl. Workshop on Logic in Databases, San Miniato, Italy, July 1-2, 1996. (with Charlie Ward).

42. P. Bonatti, M.-L. Sapino and V.S. Subrahmanian. (1996) *Merging Heterogeneous Security Orderings*, Proc. European Symposium on Research in Computer Security (ESORICS), Sep. 1996, Springer LNCS Vol. 1146, pps 183–197.

43. K.S.Candan, B. Prabhakaran and V.S. Subrahmanian. (1996) *CHIMP: A Framework for Supporting Multimedia Document Authoring and Presentation*, accepted for publication in: Proc. 1996 ACM Multimedia 1996 Conference, Boston, MA, Nov. 1996.

44. E. Bertino, S. Jajodia, P. Samarati and V.S. Subrahmanian. (1997) *A Unified Framework for Enforcing Multiple Access Control Policies*, accepted for publication in: *Proc. 1997 ACM SIGMOD Conf. on Management of Data*, Tucson, Arizona. May 1997.

45. S. Jajodia, P. Samarati and V.S. Subrahmanian. (1997) *A Logical Language for Expressing Authorizations*, accepted for publication in: *Proc. 1997 Oakland Conf. Computer Society.*

46. A. Dekhtyar and V.S. Subrahmanian. (1997) *Hybrid Probabilistic Progrms*, accepted for publication in: *Proc. · 1997 Intl. Conf. on Logic Programming*, Leuven, Belgium, July 8-12, 1997/

47. T. Eiter, J. Lu and V.S. Subrahmanian. (1997) *Computing Non-Ground representations of Stable Models*, accepted for publication in: *Proc. 1997 Intl. Conference on Logic Programming and Non-monotonic Reasoning*, Dagstuhl, Germany, July 1997.

48. E. Hwang, P. Prabhakaran, and V.S. Subrahmanian. *Distributed Video Presentations*, Accepted for publication in 1998 IEEE Intl. Conf. on Data Engineering.

*U.S. GOVERNMENT PRINTING OFFICE 1998-610-130-81024

DISTRIBUTION LIST

addresses	number of copies
JOSEPH A. CAROZZONI AFRL/IFTB 525 BROOKS ROAD ROME NY 13441-4505	1
UNIVERSITY OF MARYLAND OFFICE OF RESEARCH ADMINISTRATION AND ADVANCEMENT 2100 LEE BUILDING COLLEGE PARK MD 20742	1
AFRL/IFOIL TECHNICAL LIBRARY 26 ELECTRONIC PKY ROME NY 13441-4514	1
ATTENTION: DTIC-OCC DEFENSE TECHNICAL INFO CENTER 8725 JOHN J. KINGMAN ROAD, STE 0944 FT. BELVOIR, VA 22060-6218	2
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY 3701 NORTH FAIRFAX DRIVE ARLINGTON VA 22203-1714	1
SOFTWARE ENGR'G INST TECH LIBRARY ATTN: MR DENNIS SMITH CARNEGIE MELLON UNIVERSITY PITTSBURGH PA 15213-3890	1
USC-IST ATTN: DR ROBERT M. BALZER 4676 ADMIRALTY WAY MARINA DEL REY CA 90292-6695	1
KESTREL INSTITUTE ATTN: DR CORDELL GREEN 1801 PAGE MILL ROAD PALO ALTO CA 94304	1

ROCHESTER INSTITUTE OF TECHNOLOGY 1
ATTN: PROF J. A. LASKY
1 LOMB MEMORIAL DRIVE
P.O. BOX 9887
ROCHESTER NY 14613-5700

AFIT/ENG 1
ATTN:TOM HARTRUM
WPAFB OH 45433-6583

THE MITRE CORPORATION 1
ATTN: MR EDWARD H. BENSLEY
BURLINGTON RD/MAIL STOP A350
BEDFORD MA 01730

UNIV OF ILLINOIS, URBANA-CHAMPAIGN 1
ATTN: ANDREW CHIEN
DEPT OF COMPUTER SCIENCES
1304 W. SPRINGFIELD/240 DIGITAL LAB
URBANA IL 61801

HONEYWELL, INC. 1
ATTN: MR BERT HARRIS
FEDERAL SYSTEMS
7900 WESTPARK DRIVE
MCLEAN VA 22102

SOFTWARE ENGINEERING INSTITUTE 1
ATTN: MR WILLIAM E. HEFLEY
CARNEGIE-MELLON UNIVERSITY
SEI 2218
PITTSBURGH PA 15213-38990

UNIVERSITY OF SOUTHERN CALIFORNIA 1
ATTN: DR. YIGAL ARENS
INFORMATION SCIENCES INSTITUTE
4676 ADMIRALTY WAY/SUITE 1001
MARINA DEL REY CA 90292-6695

COLUMBIA UNIV/DEPT COMPUTER SCIENCE 1
ATTN: DR GAIL E. KAISER
450 COMPUTER SCIENCE BLDG
500 WEST 120TH STREET
NEW YORK NY 10027

SOFTWARE PRODUCTIVITY CONSORTIUM 1
ATTN: MR ROBERT LAI
2214 ROCK HILL ROAD
HERNDON VA 22070

AFIT/ENG 1
ATTN: DR GARY B. LAMONT
SCHOOL OF ENGINEERING
DEPT ELECTRICAL & COMPUTER ENGRG
WPAFB OH 45433-6583

NSA/OFC OF RESEARCH 1
ATTN: MS MARY ANNE OVERMAN
9800 SAVAGE ROAD
FT GEORGE G. MEADE MD 20755-6000

AT&T BELL LABORATORIES 1
ATTN: MR PETER G. SELFRIDGE
ROOM 3C-441
600 MOUNTAIN AVE
MURRAY HILL NJ 07974

ODYSSEY RESEARCH ASSOCIATES, INC. 1
ATTN: MS MAUREEN STILLMAN
301A HARRIS B. DATES DRIVE
ITHACA NY 14850-1313

TEXAS INSTRUMENTS INCORPORATED 1
ATTN: DR DAVID L. WELLS
P.O. BOX 655474, MS 238
DALLAS TX 75265

TEXAS A & M UNIVERSITY 1
ATTN: DR PAULA MAYER
KNOWLEDGE BASED SYSTEMS LABORATORY
DEPT OF INDUSTRIAL ENGINEERING
COLLEGE STATION TX 77843

KESTREL DEVELOPMENT CORPORATION 1
ATTN: DR RICHARD JULLIG
3260 HILLVIEW AVENUE
PALO ALTO CA 94304

DARPA/ITO 1
ATTN: DR KIRSTIE BELLMAN
3701 N FAIRFAX DRIVE
ARLINGTON VA 22203-1714

NASA/JOHNSON SPACE CENTER 1
ATTN: CHRIS CULBERT
MAIL CODE PT4
HOUSTON TX 77058

SAIC 1
ATTN: LANCE MILLER
MS T1-6-3
PO BOX 1303 (OR 1710 GOODRIDGE DR)
MCLEAN VA 22102

STERLING IMD INC. 1
KSC OPERATIONS
ATTN: MARK MAGINN
BEECHES TECHNICAL CAMPUS/RT 26 N.
ROME NY 13440

NAVAL POSTGRADUATE SCHOOL 1
ATTN: BALA RAMESH
CODE AS/RS
ADMINISTRATIVE SCIENCES DEPT
MONTEREY CA 93943

HUGHES SPACE & COMMUNICATIONS 1
ATTN: GERRY BARKSDALE
P. O. BOX 92919
BLDG R11 MS M352
LOS ANGELES, CA 90009-2919

SCHLUMBERGER LABORATORY FOR 1
COMPUTER SCIENCE
ATTN: DR. GUILLERMO ARANGO
8311 NORTH FM620
AUSTIN, TX 78720

DECISION SYSTEMS DEPARTMENT 1
ATTN: PROF WALT SCACCHI
SCHOOL OF BUSINESS
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CA 90089-1421

SOUTHWEST RESEARCH INSTITUTE 1
ATTN: BRUCE REYNOLDS
6220 CULEBRA ROAD
SAN ANTONIO, TX 78228-0510

NATIONAL INSTITUTE OF STANDARDS 1
AND TECHNOLOGY
ATTN: CHRIS DABROWSKI
ROOM A266, BLDG 225
GAITHSBURG MD 20899

EXPERT SYSTEMS LABORATORY 1
ATTN: STEVEN H. SCHWARTZ
NYNEX SCIENCE & TECHNOLOGY
500 WESTCHESTER AVENUE
WHITE PLAINS NY 20604

NAVAL TRAINING SYSTEMS CENTER 1
ATTN: ROBERT BREAUX/CODE 252
12350 RESEARCH PARKWAY
ORLANDO FL 32826-3224

CENTER FOR EXCELLENCE IN COMPUTER- 1
AIDED SYSTEMS ENGINEERING
ATTN: PERRY ALEXANDER
2291 IRVING HILL ROAD
LAWRENCE KS 66049

DR JOHN SALASIN 1
DARPA/ITO
3701 NORTH FAIRFAX DRIVE
ARLINGTON VA 22203-1714

DR BARRY BOEHM 1
DIR, USC CENTER FOR SW ENGINEERING
COMPUTER SCIENCE DEPT
UNIV OF SOUTHERN CALIFORNIA
LOS ANGELES CA 90089-0781

DR STEVE CROSS 1
CARNEGIE MELLON UNIVERSITY
SCHOOL OF COMPUTER SCIENCE
PITTSBURGH PA 15213-3891

DR MARK MAYBURY 1
MITRE CORPORATION
ADVANCED INFO SYS TECH; G041
BURLINTON ROAD, M/S K-329
BEDFORD MA 01730

ISX 1
ATTN: MR. SCOTT FOUSE
4353 PARK TERRACE DRIVE
WESTLAKE VILLAGE,CA 91361

MR GARY EDWARDS 1
ISX
433 PARK TERRACE DRIVE
WESTLAKE VILLAGE CA 91361

DR ED WALKER 1
BBN SYSTEMS & TECH CORPORATION
10 MOULTON STREET
CAMBRIDGE MA 02238

LEE ERMAN 1
CIMFLEX TEKNOWLEDGE
1810 EMBACADERO ROAD
P.O. BOX 10119
PALO ALTO CA 94303

DR. DAVE GUNNING 1
DARPA/ISO
3701 NORTH FAIRFAX DRIVE
ARLINGTON VA 22203-1714

DAN WELD 1
UNIVERSITY OF WASHINGTON
DEPART OF COMPUTER SCIENCE & ENGIN
BOX 352350
SEATTLE, WA 98195-2350

STEPHEN SODERLAND 1
UNIVERSITY OF WASHINGTON
DEPT OF COMPUTER SCIENCE & ENGIN
BOX 352350
SEATTLE, WA 98195-2350

DR. MICHAEL PITTARELLI 1
COMPUTER SCIENCE DEPART
SUNY INST OF TECH AT UTICA/ROME
P.O. BOX 3050
UTICA, NY 13504-3050

CAPRARO TECHNOLOGIES, INC 1
ATTN: GERARD CAPRARO
311 TURNER ST.
UTICA, NY 13501

USC/ISI 1
ATTN: BOB MCGREGOR
4676 ADMIRALTY WAY
MARINA DEL REY, CA 90292

SRI INTERNATIONAL 1
ATTN: ENRIQUE RUSPINI
333 RAVENSWOOD AVE
MENLO PARK, CA 94025

DARTMOUTH COLLEGE 1
ATTN: DANIELA RUS
DEPT OF COMPUTER SCIENCE
11 ROPE FERRY ROAD
HANOVER, NH 03755-3510

UNIVERSITY OF FLORIDA
ATTN: ERIC HANSON
CISE DEPT 456 CSE
GAINESVILLE, FL 32611-6120 1

CARNEGIE MELLON UNIVERSITY
ATTN: TOM MITCHELL
COMPUTER SCIENCE DEPARTMENT
PITTSBURGH, PA 15213-3890 1

CARNEGIE MELLON UNIVERSITY
ATTN: MARK CRAVEN
COMPUTER SCIENCE DEPARTMENT
PITTSBURGH, PA 15213-3890 1

UNIVERSITY OF ROCHESTER
ATTN: JAMES ALLEN
DEPARTMENT OF COMPUTER SCIENCE
ROCHESTER, NY 14627 1

TEXTWISE, LLC
ATTN: LIZ LIDDY
2-121 CENTER FOR SCIENCE & TECH
SYRACUSE, NY 13244 1

WRIGHT STATE UNIVERSITY
ATTN: DR. BRUCE BERRA
DEPART OF COMPUTER SCIENCE & ENGIN
DAYTON, OHIO 45435-0001 1

UNIVERSITY OF FLORIDA
ATTN: SHARMA CHAKRAVARTHY
COMPUTER & INFOR SCIENCE DEPART
GAINESVILLE, FL 32622-6125 1

KESTREL INSTITUTE
ATTN: DAVID ESPINOSA
3260 HILLVIEW AVENUE
PALO ALTO, CA 94304 1

STOLLER-HENKE ASSOCIATES
ATTN: T.J. GOAN
2016 BELLE MONTI AVENUE
BELMONT, CA 94002 1

USC/INFORMATION SCIENCE INSTITUTE 1
ATTN: DR. CARL KESSELMAN
11474 ADMIRALTY WAY, SUITE 1001
MARINA DEL REY, CA 90292

MASSACHUSETTS INSTITUTE OF TECH 1
ATTN: DR. MICHAELE SIEGEL
SLOAN SCHOOL
77 MASSACHUSETTS AVENUE
CAMBRIDGE, MA 02139

USC/INFORMATION SCIENCE INSTITUTE 1
ATTN: DR. WILLIAM SWARTHOUT
11474 ADMIRALTY WAY, SUITE 1001
MARINA DEL REY, CA 90292

STANFORD UNIVERSITY 1
ATTN: DR. GIO WIEDERHOLD
857 SIERRA STREET
STANFORD
SANTA CLARA COUNTY, CA 94305-4125

NCCOSC RDTE DIV D44208 1
ATTN: LEAH WONG
53245 PATTERSON ROAD
SAN DIEGO, CA 92152-7151

SPAWAR SYSTEM CENTER 1
ATTN: LES ANDERSON
271 CATALINA BLVD, CODE 413
SAN DIEGO CA 92151

GEORGE MASON UNIVERSITY 1
ATTN: SUSHIL JAJODIA
ISSE DEPT
FAIRFAX, VA 22030-4444

DIRNSA 1
ATTN: MICHAEL R. WARE
DOD, NSA/CSS (R23)
FT. GEORGE G. MEADE MD 20755-6000

DR. JIM RICHARDSON 1
3660 TECHNOLOGY DRIVE
MINNEAPOLIS, MN 55418

LOUISIANA STATE UNIVERSITY COMPUTER SCIENCE DEPT ATTN: DR. PETER CHEN 257 COATES HALL BATON ROUGE, LA 70803	1
INSTITUTE OF TECH DEPT OF COMP SCI ATTN: DR. JAIDEEP SRIVASTAVA 4-192 EE/CS 200 UNION ST SE MINNEAPOLIS, MN 55455	1
GTE/BBN ATTN: MAURICE M. MCNEIL 9655 GRANITE RIDGE DRIVE SUITE 245 SAN DIEGO, CA 92123	1
UNIVERSITY OF FLORIDA ATTN: DR. SHARMA CHAKRAVARTHY E470 CSE BUILDING GAINESVILLE, FL 32611-6125	1

MISSION
OF
AFRL/INFORMATION DIRECTORATE (IF)

The advancement and application of information systems science and technology for aerospace command and control and its transition to air, space, and ground systems to meet customer needs in the areas of Global Awareness, Dynamic Planning and Execution, and Global Information Exchange is the focus of this AFRL organization. The directorate's areas of investigation include a broad spectrum of information and fusion, communication, collaborative environment and modeling and simulation, defensive information warfare, and intelligent information systems technologies.